

LISTING OF CLAIMS

This listing of claims will replace all prior versions and listings of claims in the application:

1. (Currently Amended) A method for efficiently batching transactions in a computerized system, comprising the steps of:
 - receiving a transaction at the computerized system;
 - writing the received transaction data to a data buffer;
 - if the received transaction is an only transaction in the computerized system, then flushing the data buffer to a stable non-volatile storage device;
 - if the received transaction is not the only transaction in the computerized system, then determining whether a maximum response time will be met for an oldest transaction waiting for flush in the computerized system if the data buffer were flushed to the stable storage device immediately;
 - if the maximum response time will not be met, then flushing the data buffer to the stable non-volatile storage device;
 - if the maximum response time will be met, then determining whether a rate of data of incoming transactions is higher than a rate of flushing data from the data buffer to the stable non-volatile storage device if flushed immediately;
 - if the rate of data of incoming transactions is lower than the rate of flushing data from the data buffer to the stable non-volatile storage device if flushed immediately, then flushing the data buffer to the stable non-volatile storage device;
 - if the rate of incoming data is higher than the rate of flushing data from the data buffer to the stable non-volatile storage device if flushed immediately, then waiting to receive another transaction; and
 - updating a flush performance history after each flush of the data buffer to the stable non-volatile storage device.
2. (Currently Amended) The method of claim 1 wherein the step of determining whether the maximum response time will be met, comprises the steps of:

determining a time queued of the oldest transaction waiting for flush in the computerized system;

determining an amount of data in the data buffer;

using the flush performance history to compute a time duration to flush the data buffer to the stable non-volatile storage device based on the amount of data in the data buffer;

calculating an expected delay in flushing the oldest transaction to the stable non-volatile storage device by subtracting the time when the oldest transaction announced that it was waiting for flush from a current time and adding the time to flush the data buffer to the stable non-volatile storage device;

comparing the expected delay to a maximum response value;

if the expected delay is not less than the maximum response value, then flushing the data buffer to the stable non-volatile storage device; and

if the expected delay is less than the maximum response value, then proceeding to the step of determining the rate of incoming data.

3. (Currently Amended) The method of claim 1 wherein the step of determining whether a rate of incoming data of incoming transactions associated with the received transaction is higher than a the current rate of flushing data from the data buffer to the stable non-volatile storage device comprises the steps of:

determining a recent rate of incoming transactions;

determining an amount of data submitted by an average transaction to the data buffer by maintaining amounts of data for a specified number of recent transactions and averaging the maintained amounts of data;

computing a rate per second of an amount of data submitted to the data buffer by the received incoming transactions;

determining a target amount of data in the data buffer that is needed to achieve sufficient batching in a flush operation based on the rate of data of incoming transactions data and the flush performance history;

determining an amount of data in the data buffer;

comparing the amount of data in the data buffer to the target amount of data; and

if the amount of data in the data buffer is not greater than the target amount of data, then waiting for a next received transaction; and

if the amount of data in the data buffer is greater than the target amount of data, then flushing the data buffer to the stable non-volatile storage device.

4. (Currently Amended) The method of claim 1 wherein the step of flushing the data buffer to the stable non-volatile storage device comprises the steps of:

determining how many transactions are waiting for a flush of the data buffer;

determining an amount of data in the data buffer;

persisting the data in the data buffer to the stable non-volatile storage device;

marking previously used memory by the data buffer as available for re-use;

resetting a data buffer size parameter to zero;

determining a duration for the writing, marking and resetting steps;

updating the data buffer flush performance history with the determined duration for writing, marking and resetting;

determining an amount of data submitted by an average transaction to the data buffer by maintaining amounts of data for a specified number of recent transactions and averaging the maintained amounts of data;

updating an average transaction data buffer size parameters; and

notifying waiting transactions that the flush has completed.

5. (Currently Amended) The method of claim 1 wherein the flush performance history is a table having a plurality of buckets, each bucket having a minimum and a maximum amount of data buffer data for which that bucket holds statistics, each bucket capable of storing a plurality of entries, wherein each entry is an observed time to flush the data in the data buffer to the stable non-volatile storage device, wherein the step of updating the table of rates comprises the steps of:

selecting a bucket in the plurality in which to record a data update based on the amount of data in the data buffer; and

entering the data update into the selected bucket wherein the data update includes a time duration for flushing the data in the data buffer to the stable non-volatile storage device.

6. (Original) The method of claim 5 wherein the step of updating the table of rates further comprises the steps of:

comparing the number of entries in the selected bucket to a bucket history length constant; and

if the number of entries in the selected bucket is more than the bucket history length constant, then removing the oldest entry.

7. (Currently Amended) The method of claim 3 wherein the flush performance history is a table having a plurality of buckets, each bucket having a minimum and a maximum amount of data buffer data for which that bucket holds statistics, each bucket capable of storing a plurality of entries, wherein each entry is an observed time to flush the data buffer to the stable non-volatile storage device and wherein the step of determining a target amount of data comprises the steps of:

- a) selecting a first bucket in the flush performance history;
- b) determining a minimum and maximum amount of data buffer data for which the selected bucket holds flush statistics ;
- c) based on the minimum and maximum amount of data, determining an average amount of data in the data buffer for which the selected bucket holds flush performance data;
- d) calculating an average duration of a flush of the data buffer by averaging values in the plurality of entries in the selected bucket;
- e) calculating a number of times per second that a flush can be called based on the average duration of a flush for the selected bucket;
- f) calculating a rate of flushing for the selected bucket;

- g) comparing the rate of flushing for the selected bucket to the rate of incoming data;
- h) if the rate of flushing for the selected bucket is not less than the rate of incoming data, then taking the average data amount of data in the data buffer for the selected bucket as the target amount of data;
- i) if the rate of flushing for the selected bucket is less than the rate of incoming data, determining whether there is at least one other bucket in the flush performance history;
- j) if there is at least one other bucket in the flush performance history, selecting a second bucket in the flush performance history and repeating steps b) - j); and
- k) if there is not at least one other bucket in the flush performance history, then taking the average data amount of data in the data buffer for the selected bucket as the target amount of data.

8. (Original) A method for efficiently batching transactions in a database system, comprising the steps of:

- receiving a transaction at the database system;
- writing the received transaction data to a data buffer;
- if the received transaction is an only transaction in the computerized system, then flushing the data buffer to a database;
- if the received transaction is not the only transaction in the computerized system, then determining whether a maximum response time will be met for an oldest transaction waiting for flush in the computerized system if the data buffer were flushed to the database immediately;
- if the maximum response time will not be met, then flushing the data buffer to the database;
- if the maximum response time will be met, then determining whether a rate of data of incoming transactions is higher than a rate of flushing data from the data buffer to the database if flushed immediately;

if the rate of data of incoming transactions is lower than the rate of flushing data from the data buffer to the database if flushed immediately, then flushing the data buffer to the database;

if the rate of incoming data is higher than the rate of flushing data from the data buffer to the database if flushed immediately, then waiting to receive another transaction; and

updating a flush performance history after each flush of the data buffer to the database.

9. (Currently Amended) A transaction system that efficiently batches transactions, comprising:

a data interface to receive transactions;

a stable non-volatile storage device to store processed transactions;

a memory including a data buffer and a system performance history, the system performance history including system performance data about rates of flushing the data buffer to the stable non-volatile storage device and batch sizes of flushing the data buffer to the stable non-volatile storage device; and

a controller coupled to the interface, the stable non-volatile storage device and the memory, the controller configured to process the transactions, to write the processed transactions to the data buffer, to flush the data buffer to the stable non-volatile storage device according to the system performance data from the system performance history, and to update the system performance history with data from the flush of the data buffer to the stable non-volatile storage device.

10. (Currently Amended) The system of claim 9 wherein the controller further comprises:

a transaction processor to process transaction and to write the processed transactions; and

a flush processor to flush the data buffer to the stable non-volatile storage device and to update the system performance history.

11. (Currently Amended) The system of claim 9 wherein the controller further comprises a maximum response time value, the controller being further configured to determine whether a maximum response time will be met for an oldest transaction in the data buffer and if the maximum response time will not be met, to flush the data buffer to the stable non-volatile storage device.
12. (Original) The system of claim 11 wherein the response time is determined based on information from the system performance history.
13. (Original) The system of claim 9 wherein the controller is configured to determine a rate of incoming data and to determine a rate of flushing data, and if the rate of incoming data is higher than the rate of flushing data, then to wait for further transactions in order to develop a larger batch of processed transactions.
14. (Currently Amended) The system of claim 9 wherein the controller is further configured to adjust frequency of flushing and amount of data per flush to the stable non-volatile storage device based on information from the system performance history.
15. (Currently Amended) The system of claim 9 wherein the system performance history is a table having a plurality of buckets, each bucket having a minimum and a maximum amount of data buffer data for which that bucket holds statistics, each bucket capable of storing a plurality of entries, and wherein each entry is an observed time to flush the data in the data buffer to the stable non-volatile storage device.
16. (Currently Amended) The system of claim 15 wherein the controller is further configured to select a bucket in the plurality of buckets in the system performance history into which to record a data update, the controller to select based on an amount of data in the data buffer, the controller is further configured to enter the data update

into the selected bucket where the data update includes a time duration for flushing data in the data buffer to the stable non-volatile storage device.

17. (Original) The system of claim 15 wherein the system performance history includes a bucket history length constant, and wherein the controller is configured to compare the number of entries in the selected bucket to the bucket history length constant and if the number of entries exceeds the bucket history length constant, to remove old entries in excess of the bucket history length constant.

18. (Currently Amended) A database system that efficiently batches transactions, comprising:

a data interface to receive database transactions;

a stable non-volatile storage device having a database to store processed database transactions;

a memory including a data buffer and a system performance history, the system performance history including system performance data about rates of flushing the data buffer to the database and batch sizes of flushing the data buffer to the database; and

a controller coupled to the interface, the stable non-volatile storage device and the memory, the controller configured to process the transactions, to write the processed transactions to the data buffer, to flush the data buffer to the database according to the system performance data from the system performance history, and to update the system performance history with data from the flush of the data buffer to the database.

19. (Currently Amended) A computer program product having a computer-readable medium including computer program logic encoded thereon that, when performed on a computer system directs the computer system to perform the method of:

receiving a transaction at the computerized system;

writing the received transaction data to a data buffer;

if the received transaction is an only transaction in the computerized system, then flushing the data buffer to a stable non-volatile storage device;

if the received transaction is not the only transaction in the computerized system, then determining, in the computerized system, whether a maximum response time will be met for an oldest transaction waiting for flush in the computerized system if the data buffer were flushed to the stable non-volatile storage device immediately;

if the maximum response time will not be met, then flushing the data buffer to the stable non-volatile storage device;

if the maximum response time will be met, then determining, in the computerized system, whether a rate of data of incoming transactions is higher than a rate of flushing data from the data buffer to the stable non-volatile storage device if flushed immediately;

if the rate of data of incoming transactions is lower than the rate of flushing data from the data buffer to the stable non-volatile storage device if flushed immediately, then flushing the data buffer to the stable non-volatile storage device;

if the rate of incoming data is higher than the rate of flushing data from the data buffer to the stable non-volatile storage device if flushed immediately, then waiting to receive another transaction; and

updating a flush performance history in the computerized system after each flush of the data buffer to the stable non-volatile storage device.

20. (Currently Amended) A system for efficiently batching transactions in a computerized system, comprising:

means for receiving a transaction at the computerized system;

means for writing the received transaction data to a data buffer;

means for flushing the data buffer to a stable non-volatile storage device if the received transaction is an only transaction in the computerized system;

means for determining whether a maximum response time will be met for an oldest transaction waiting for flush in the computerized system if the data buffer were

flushed to the stable non-volatile storage device immediately if the received transaction is not the only transaction in the computerized system;

means for flushing the data buffer to the stable non-volatile storage device if the maximum response time will not be met;

means for determining whether a rate of data of incoming transactions is higher than a rate of flushing data from the data buffer to the stable non-volatile storage device if flushed immediately if the maximum response time will be met;

means for flushing the data buffer to the stable non-volatile storage device if the rate of data of incoming transactions is lower than the rate of flushing data from the data buffer to the stable non-volatile storage device if flushed immediately;

means for waiting to receive another transaction if the rate of incoming data is higher than the rate of flushing data from the data buffer to the stable non-volatile storage device if flushed immediately; and

means for updating a flush performance history after each flush of the data buffer to the stable non-volatile storage device.